

Internet Engineering Task Force (IETF)
Request for Comments: 8779
Category: Standards Track
ISSN: 2070-1721

C. Margaria, Ed.
Juniper
O. Gonzalez de Dios, Ed.
Telefonica Investigacion y Desarrollo
F. Zhang, Ed.
Huawei Technologies
July 2020

Path Computation Element Communication Protocol (PCEP) Extensions for GMPLS

Abstract

A Path Computation Element (PCE) provides path computation functions for Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks. Additional requirements for GMPLS are identified in RFC 7025.

This memo provides extensions to the Path Computation Element Communication Protocol (PCEP) for the support of the GMPLS control plane to address those requirements.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8779>.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
 - 1.1. Terminology
 - 1.2. PCEP Requirements for GMPLS
 - 1.3. Requirements Applicability
 - 1.3.1. Requirements on the Path Computation Request
 - 1.3.2. Requirements on the Path Computation Response
 - 1.4. Existing Support and Limitations for GMPLS in Base PCEP Objects

- 2. PCEP Objects and Extensions
 - 2.1. GMPLS Capability Advertisement
 - 2.1.1. GMPLS Computation TLV in the Existing PCE Discovery Protocol
 - 2.1.2. OPEN Object Extension GMPLS-CAPABILITY TLV
 - 2.2. RP Object Extension
 - 2.3. BANDWIDTH Object Extensions
 - 2.4. LOAD-BALANCING Object Extensions
 - 2.5. END-POINTS Object Extensions
 - 2.5.1. Generalized Endpoint Object Type
 - 2.5.2. END-POINTS TLV Extensions
 - 2.6. IRO Extension
 - 2.7. XRO Extension
 - 2.8. LSPA Extensions
 - 2.9. NO-PATH Object Extension
 - 2.9.1. Extensions to NO-PATH-VECTOR TLV
 - 3. Additional Error-Types and Error-Values Defined
 - 4. Manageability Considerations
 - 4.1. Control of Function through Configuration and Policy
 - 4.2. Information and Data Models
 - 4.3. Liveness Detection and Monitoring
 - 4.4. Verifying Correct Operation
 - 4.5. Requirements on Other Protocols and Functional Components
 - 4.6. Impact on Network Operation
 - 5. IANA Considerations
 - 5.1. PCEP Objects
 - 5.2. Endpoint Type Field in the Generalized END-POINTS Object
 - 5.3. New PCEP TLVs
 - 5.4. RP Object Flag Field
 - 5.5. New PCEP Error Codes
 - 5.6. New Bits in NO-PATH-VECTOR TLV
 - 5.7. New Subobject for the Include Route Object
 - 5.8. New Subobject for the Exclude Route Object
 - 5.9. New GMPLS-CAPABILITY TLV Flag Field
 - 6. Security Considerations
 - 7. References
 - 7.1. Normative References
 - 7.2. Informative References
- Appendix A. LOAD-BALANCING Usage for SDH Virtual Concatenation
- Acknowledgments
- Contributors
- Authors' Addresses

1. Introduction

Although the PCE architecture and framework for both MPLS and GMPLS networks are defined in [RFC4655], most pre-existing PCEP RFCs, such as [RFC5440], [RFC5521], [RFC5541], and [RFC5520], are focused on MPLS networks and do not cover the wide range of GMPLS networks. This document complements these RFCs by addressing the extensions required for GMPLS applications and routing requests, for example, for Optical Transport Networks (OTNs) and Wavelength Switched Optical Networks (WSOs).

The functional requirements to be addressed by the PCEP extensions to support these applications are fully described in [RFC7025] and [RFC7449].

1.1. Terminology

This document uses terminologies from the PCE architecture document [RFC4655]; the PCEP documents including [RFC5440], [RFC5521], [RFC5541], [RFC5520], [RFC7025], and [RFC7449]; and the GMPLS documents such as [RFC3471], [RFC3473], and so on. Note that the reader is expected to be familiar with these documents. The following abbreviations are used in this document:

ERO: Explicit Route Object
 IRO: Include Route Object
 L2SC: Layer 2 Switch Capable [RFC3471]
 LSC: Lambda Switch Capable [RFC3471]
 LSP: Label Switched Path
 LSPA: LSP Attribute
 MEF: Metro Ethernet Forum
 MT: Multiplier [RFC4328] [RFC4606]
 NCC: Number of Contiguous Components [RFC4606]
 NVC: Number of Virtual Components [RFC4328] [RFC4606]
 ODU: Optical Data Unit [G.709-v3]
 OTN: Optical Transport Network [G.709-v3]
 P2MP: Point-to-Multipoint
 PCC: Path Computation Client
 PCRep: Path Computation Reply [RFC5440]
 PCReq: Path Computation Request [RFC5440]
 RCC: Requested Contiguous Concatenation [RFC4606]
 RRO: Record Route Object
 RSVP-TE: Resource Reservation Protocol - Traffic Engineering
 SDH: Synchronous Digital Hierarchy
 SONET: Synchronous Optical Network
 SRLG: Shared Risk Link Group
 SSON: Spectrum-Switched Optical Network
 TDM: Time-Division Multiplex Capable [RFC3471]
 TE-LSP: Traffic Engineered LSP
 XRO: Exclude Route Object

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. PCEP Requirements for GMPLS

[RFC7025] describes the set of PCEP requirements that support GMPLS TE-LSPs. This document assumes a significant familiarity with [RFC7025] and existing PCEP extensions. As a short overview, those requirements can be broken down into the following categories.

- * Which data flow is switched by the LSP: a combination of a switching type (for instance, L2SC or TDM), an LSP encoding type (e.g., Ethernet, SONET/SDH), and sometimes the signal type (e.g., in case of a TDM or an LSC switching capability).
- * Data-flow-specific traffic parameters, which are technology specific. For instance, in SDH/SONET and OTN networks [G.709-v3], the concatenation type and the concatenation number have an influence on the switched data and on which link it can be supported.
- * Support for asymmetric bandwidth requests.
- * Support for unnumbered interface identifiers, as defined in [RFC3477].
- * Label information and technology-specific label(s) such as wavelength labels as defined in [RFC6205]. A PCC should also be able to specify a label restriction similar to the one supported by RSVP-TE in [RFC3473].
- * Ability to indicate the requested granularity for the path ERO: node, link, or label. This is to allow the use of the explicit label control feature of RSVP-TE.

The requirements of [RFC7025] apply to several objects conveyed by PCEP; this is described in Section 1.3. Some of the requirements of [RFC7025] are already supported in existing documents, as described in Section 1.4.

This document describes a set of PCEP extensions, including new object types, TLVs, encodings, error codes, and procedures, in order to fulfill the aforementioned requirements not covered in existing RFCs.

1.3. Requirements Applicability

This section follows the organization of [RFC7025], Section 3 and indicates, for each requirement, the affected piece of information carried by PCEP and its scope.

1.3.1. Requirements on the Path Computation Request

- (1) Switching capability/type: As described in [RFC3471], this piece of information is used with the encoding type and signal type to fully describe the switching technology and data carried by the TE-LSP. This is applicable to the TE-LSP itself and also to the TE-LSP endpoint (carried in the END-POINTS object for MPLS networks in [RFC5440]) when considering multiple network layers. Inter-layer path computation requirements are addressed in [RFC8282], which focuses on the TE-LSP itself but does not address the TE-LSP endpoints.
- (2) Encoding type: See (1).
- (3) Signal type: See (1).
- (4) Concatenation type: This parameter and the concatenation number (see (5)) are specific to some TDM (SDH and ODU) switching technologies. They MUST be described together and are used to derive the requested resource allocation for the TE-LSP. It is scoped to the TE-LSP and is related to the BANDWIDTH object [RFC5440] in MPLS networks. See concatenation information in [RFC4606] and [RFC4328].
- (5) Concatenation number: See (4).

- (6) Technology-specific label(s): As described in [RFC3471], the GMPLS labels are specific to each switching technology. They can be specified on each link and also on the TE-LSP endpoints, in WSON networks, for instance, as described in [RFC6163]. The label restriction can apply to endpoints, and on each hop, the related PCEP objects are END-POINTS, IRO, XRO, and RRO.
- (7) End-to-End (E2E) path protection type: As defined in [RFC4872], this is applicable to the TE-LSP. In MPLS networks, the related PCEP object is LSPA (carrying local protection information).
- (8) Administrative group: As defined in [RFC3630], this information is already carried in the LSPA object.
- (9) Link protection type: As defined in [RFC4872], this is applicable to the TE-LSP and is carried in association with the E2E path protection type.
- (10) Support for unnumbered interfaces: As defined in [RFC3477]. Its scope and related objects are the same as labels.
- (11) Support for asymmetric bandwidth requests: As defined in [RFC6387], the scope is similar to (4).
- (12) Support for explicit label control during the path computation: This affects the TE-LSP and the amount of information returned in the ERO.
- (13) Support of label restrictions in the requests/responses: This is described in (6).

1.3.2. Requirements on the Path Computation Response

- (1) Path computation with concatenation: This is related to the Path Computation request requirement (4). In addition, there is a specific type of concatenation, called virtual concatenation, that allows different routes to be used between the endpoints. It is similar to the semantic and scope of the LOAD-BALANCING in MPLS networks.
- (2) Label constraint: The PCE should be able to include labels in the path returned to the PCC; the related object is the ERO object.
- (3) Roles of the routes: As defined in [RFC4872], this is applicable to the TE-LSP and is carried in association with the E2E path protection type.

1.4. Existing Support and Limitations for GMPLS in Base PCEP Objects

The support provided by specifications in [RFC8282] and [RFC5440] for the requirements listed in [RFC7025] is summarized in Tables 1 and 2. In some cases, the support may not be complete, as noted, and additional support needs to be provided as indicated in this specification.

+=====+=====+=====+		+=====+=====+=====+	
Req.	Name	Support	
+=====+=====+=====+		+=====+=====+=====+	
1	Switching capability/type	SWITCH-LAYER	
		(RFC 8282)	
+-----+-----+-----+		+-----+-----+-----+	
2	Encoding type	SWITCH-LAYER	
		(RFC 8282)	

3	Signal type	SWITCH-LAYER (RFC 8282)
4	Concatenation type	No
5	Concatenation number	No
6	Technology-specific label	(Partial) ERO (RFC 5440)
7	End-to-End (E2E) path protection type	No
8	Administrative group	LSPA (RFC 5440)
9	Link protection type	No
10	Support for unnumbered interfaces	(Partial) ERO (RFC 5440)
11	Support for asymmetric bandwidth requests	No
12	Support for explicit label control during the path computation	No
13	Support of label restrictions in the requests/responses	No

Table 1: Requirements Support per RFC 7025, Section 3.1

Req.	Name	Support
1	Path computation with concatenation	No
2	Label constraint	No
3	Roles of the routes	No

Table 2: Requirements Support per RFC 7025, Section 3.2

Per Section 1.3, PCEP (as described in [RFC5440], [RFC5521], and [RFC8282]) supports the following objects, included in requests and responses, that are related to the described requirements.

From [RFC5440]:

END-POINTS: related to requirements 1, 2, 3, 6, 10, and 13. The object only supports numbered endpoints. The context specifies whether they are node identifiers or numbered interfaces.

BANDWIDTH: related to requirements 4, 5, and 11. The data rate is encoded in the BANDWIDTH object (as an IEEE 32-bit float). [RFC5440] does not include the ability to convey an encoding proper to all GMPLS-controlled networks.

ERO: related to requirements 6, 10, 12, and 13. The ERO content is defined in RSVP in [RFC3209], [RFC3473], [RFC3477], and [RFC7570] and already supports all of the requirements.

LSPA: related to requirements 7, 8, and 9. Requirement 8

(Administrative group) is already supported.

From [RFC5521]:

XRO:

- This object allows excluding (strict or not) resources and is related to requirements 6, 10, and 13. It also includes the requested diversity (node, link, or SRLG).
- When the F bit is set, the request indicates that the existing path has failed, and the resources present in the RRO can be reused.

From [RFC8282]:

SWITCH-LAYER: addresses requirements 1, 2, and 3 for the TE-LSP and indicates which layer(s) should be considered. The object can be used to represent the RSVP-TE Generalized Label Request. It does not address the endpoints case of requirements 1, 2, and 3.

REQ-ADAP-CAP: indicates the adaptation capabilities requested; it can also be used for the endpoints in case of mono-layer computation.

The gaps in functional coverage of the base PCEP objects are:

- * The BANDWIDTH and LOAD-BALANCING objects do not describe the details of the traffic request (requirements 4 and 5, for example, NVC and multiplier) in the context of GMPLS networks, for instance, in TDM or OTN networks.
- * The END-POINTS object does not allow specifying an unnumbered interface, nor potential label restrictions on the interface (requirements 6, 10, and 13). Those parameters are of interest in case of switching constraints.
- * The IROs/XROs do not allow the inclusion/exclusion of labels (requirements 6, 10, and 13).
- * Base attributes do not allow expressing the requested link protection level and/or the end-to-end protection attributes.

As defined later in this document, the PCEP extensions that cover the gaps are:

- * Two new object types are defined for the BANDWIDTH object (Generalized bandwidth and Generalized bandwidth of an existing TE-LSP for which a reoptimization is requested).
- * A new object type is defined for the LOAD-BALANCING object (Generalized Load Balancing).
- * A new object type is defined for the END-POINTS object (Generalized Endpoint).
- * A new TLV is added to the Open message for capability negotiation.
- * A new TLV is added to the LSPA object.
- * The Label subobject is now allowed in the IRO and XRO objects.
- * In order to indicate the routing granularity used in the response, a new flag is added in the RP object.

2. PCEP Objects and Extensions

This section describes the necessary PCEP objects and extensions. The PCReq and PCRep messages are defined in [RFC5440]. This document does not change the existing grammar.

2.1. GMPLS Capability Advertisement

2.1.1. GMPLS Computation TLV in the Existing PCE Discovery Protocol

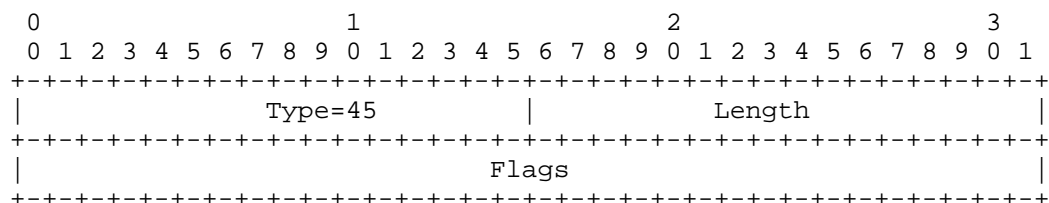
IGP-based PCE Discovery (PCED) is defined in [RFC5088] and [RFC5089] for the OSPF and IS-IS protocols. Those documents have defined bit 0 in the PCE-CAP-FLAGS Sub-TLV of the PCED TLV as "Path computation with GMPLS link constraints". This capability is optional and can be used to detect GMPLS-capable PCEs. PCEs that set the bit to indicate support of GMPLS path computation MUST follow the procedures in Section 2.1.2 to further qualify the level of support during PCEP session establishment.

2.1.2. OPEN Object Extension GMPLS-CAPABILITY TLV

In addition to the IGP advertisement, a PCEP speaker MUST be able to discover the other peer GMPLS capabilities during the Open message exchange. This capability is also useful to avoid misconfigurations. This document defines a GMPLS-CAPABILITY TLV for use in the OPEN object to negotiate the GMPLS capability. The inclusion of this TLV in the Open message indicates that the PCEP speaker supports the PCEP extensions defined in the document. A PCEP speaker that is able to support the GMPLS extensions defined in this document MUST include the GMPLS-CAPABILITY TLV in the Open message. If one of the PCEP peers does not include the GMPLS-CAPABILITY TLV in the Open message, the peers MUST NOT make use of the objects and TLVs defined in this document.

If the PCEP speaker supports the extensions of this specification but did not advertise the GMPLS-CAPABILITY capability, upon receipt of a message from the PCE including an extension defined in this document, it MUST generate a PCEP Error (PCErr) with Error-Type=10 (Reception of an invalid object) and Error-value=31 (Missing GMPLS-CAPABILITY TLV), and it SHOULD terminate the PCEP session.

As documented in Section 5.3 ("New PCEP TLVs"), IANA has allocated value 45 (GMPLS-CAPABILITY) from the "PCEP TLV Type Indicators" sub-registry. The format for the GMPLS-CAPABILITY TLV is shown in the following figure.



No flags are defined in this document; they are reserved for future use. Unassigned flags MUST be set to zero on transmission and MUST be ignored on receipt.

2.2. RP Object Extension

Explicit Label Control (ELC) is a procedure supported by RSVP-TE, where the outgoing labels are encoded in the ERO. As a consequence, the PCE can provide such labels directly in the path ERO. Depending on the policies or switching layer, it might be necessary for the PCC to use explicit label control or explicit link ids; thus, it needs to indicate in the PCReq which granularity it is expecting in the ERO.

This corresponds to requirement 12 in Section 3.1 of [RFC7025]. The possible granularities can be node, link, or label. The granularities are interdependent, in the sense that link granularity implies the presence of node information in the ERO; similarly, a label granularity implies that the ERO contains node, link, and label information.

A new 2-bit Routing Granularity (RG) flag (bits 15-16) is defined in the RP object. The values are defined as follows:

- 0: reserved
- 1: node
- 2: link
- 3: label

The RG flag in the RP object indicates the requested route granularity. The PCE SHOULD follow this granularity and MAY return a NO-PATH if the requested granularity cannot be provided. The PCE MAY return any granularity on the route based on its policy. The PCC can decide if the ERO is acceptable based on its content.

If a PCE honored the requested routing granularity for a request, it MUST indicate the selected routing granularity in the RP object included in the response. Otherwise, the PCE MUST use the reserved RG to leave the check of the ERO to the PCC. The RG flag is backward compatible with [RFC5440]: the value sent by an implementation (PCC or PCE) not supporting it will indicate a reserved value.

2.3. BANDWIDTH Object Extensions

Per [RFC5440], the object carrying the requested size for the TE-LSP is the BANDWIDTH object. Object types 1 and 2 defined in [RFC5440] do not provide enough information to describe the TE-LSP bandwidth in GMPLS networks. The BANDWIDTH object encoding has to be extended to allow the object to express the bandwidth as described in [RFC7025]. RSVP-TE extensions for GMPLS provide a set of encodings that allow such representation in an unambiguous way; this is encoded in the RSVP-TE Traffic Specification (TSpec) and Flow Specification (FlowSpec) objects. This document extends the BANDWIDTH object with new object types reusing the RSVP-TE encoding.

The following possibilities are supported by the extended encoding:

- * Asymmetric bandwidth (different bandwidth in forward and reverse direction), as described in [RFC6387].
- * GMPLS (SDH/SONET, G.709, ATM, MEF, etc.) parameters.

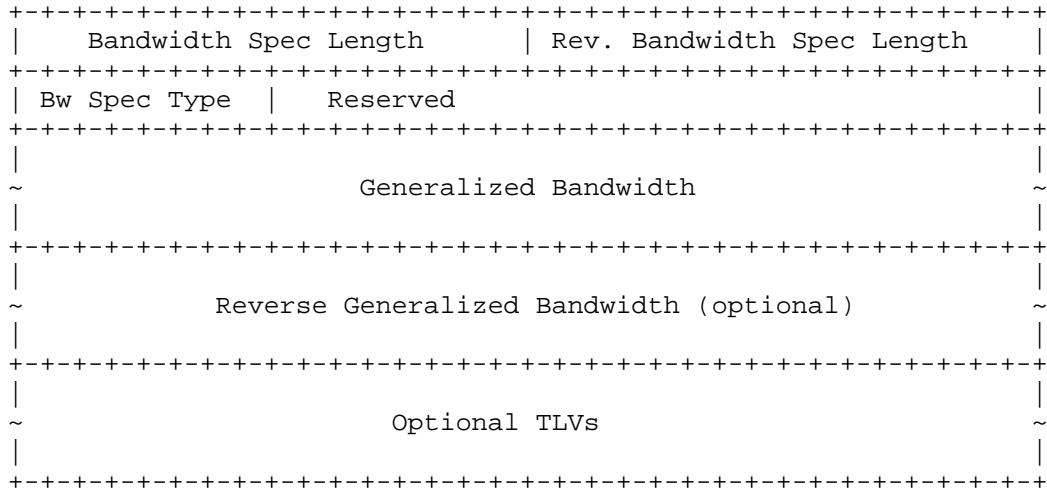
This corresponds to requirements 3, 4, 5, and 11 in Section 3.1 of [RFC7025].

This document defines two object types for the BANDWIDTH object:

- 3: Generalized bandwidth
- 4: Generalized bandwidth of an existing TE-LSP for which a reoptimization is requested

The definitions below apply for object types 3 and 4. The body is as follows:

0	1	2	3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1			



BANDWIDTH object types 3 and 4 have a variable length. The 16-bit Bandwidth Spec Length field indicates the length of the Generalized Bandwidth field. The Bandwidth Spec Length MUST be strictly greater than 0. The 16-bit Reverse Bandwidth Spec Length field indicates the length of the Reverse Generalized Bandwidth field. The Reverse Bandwidth Spec Length MAY be equal to 0.

The Bw Spec Type field determines which type of bandwidth is represented by the object.

The Bw Spec Type corresponds to the RSVP-TE SENDER_TSPEC (Object Class 12) C-Types.

The encoding of the Generalized Bandwidth and Reverse Generalized Bandwidth fields is the same as the traffic parameters carried in RSVP-TE; they can be found in the following references. Note that the RSVP-TE traffic specification MAY also include TLVs that are different from the PCEP TLVs (e.g., the TLVs defined in [RFC6003]).

Bw Spec Type	Name	Reference
2	Intserv	[RFC2210]
4	SONET/SDH	[RFC4606]
5	G.709	[RFC4328]
6	Ethernet	[RFC6003]
7	OTN-TDM	[RFC7139]
8	SSON	[RFC7792]

Table 3: Generalized Bandwidth and Reverse Generalized Bandwidth Field Encoding

When a PCC requests a bidirectional path with symmetric bandwidth, it SHOULD only specify the Generalized Bandwidth field and set the Reverse Bandwidth Spec Length to 0. When a PCC needs to request a bidirectional path with asymmetric bandwidth, it SHOULD specify the different bandwidth in the forward and reverse directions with Generalized Bandwidth and Reverse Generalized Bandwidth fields.

The procedure described in [RFC5440] for the PCRep is unchanged: a PCE MAY include the BANDWIDTH objects in the response to indicate the

BANDWIDTH of the path.

As specified in [RFC5440], in the case of the reoptimization of a TE-LSP, the bandwidth of the existing TE-LSP MUST also be included in addition to the requested bandwidth if and only if the two values differ. The object type 4 MAY be used instead of the previously specified object type 2 to indicate the existing TE-LSP bandwidth, which was originally specified with object type 3. A PCC that requested a path with a BANDWIDTH object of object type 1 MUST use object type 2 to represent the existing TE-LSP bandwidth.

Optional TLVs MAY be included within the object body to specify more specific bandwidth requirements. No TLVs for object types 3 and 4 are defined by this document.

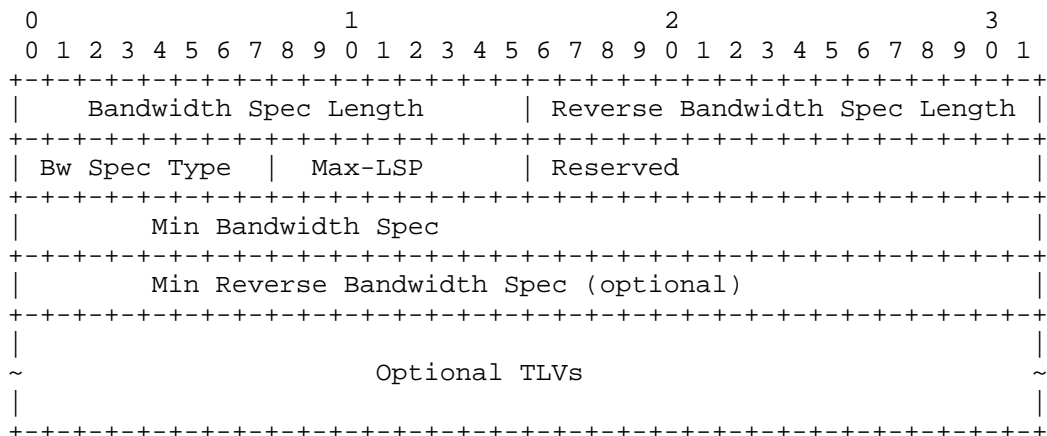
2.4. LOAD-BALANCING Object Extensions

The LOAD-BALANCING object [RFC5440] is used to request a set of at most Max-LSP TE-LSPs having in total the bandwidth specified in BANDWIDTH, with each TE-LSP having at least a specified minimum bandwidth. The LOAD-BALANCING object follows the bandwidth encoding of the BANDWIDTH object; thus, the existing definition from [RFC5440] does not describe enough details for the bandwidth specification expected by GMPLS.

Similar to the BANDWIDTH object, a new object type is defined to allow a PCC to represent the bandwidth types supported by GMPLS networks.

This document defines object type 2 (Generalized Load Balancing) for the LOAD-BALANCING object. The Generalized Load Balancing object type has a variable length.

The format of the Generalized Load Balancing object type is as follows:



Bandwidth Spec Length (16 bits): the total length of the Min Bandwidth Spec field. The length MUST be strictly greater than 0.

Reverse Bandwidth Spec Length (16 bits): the total length of the Min Reverse Bandwidth Spec field. It MAY be equal to 0.

Bw Spec Type (8 bits): the bandwidth specification type; it corresponds to RSVP-TE SENDER_TSPEC (Object Class 12) C-Types.

Max-LSP (8 bits): the maximum number of TE-LSPs in the set.

Min Bandwidth Spec (variable): specifies the minimum bandwidth specification of each element of the TE-LSP set.

Min Reverse Bandwidth Spec (variable): specifies the minimum reverse bandwidth specification of each element of the TE-LSP set.

The encoding of the Min Bandwidth Spec and Min Reverse Bandwidth Spec fields is the same as in the RSVP-TE SENDER_TSPEC object; it can be found in Table 3 in Section 2.3 of this document.

When a PCC requests a bidirectional path with symmetric bandwidth while specifying load-balancing constraints, it SHOULD specify the Min Bandwidth Spec field and set the Reverse Bandwidth Spec Length to 0. When a PCC needs to request a bidirectional path with asymmetric bandwidth while specifying load-balancing constraints, it MUST specify the different bandwidth in forward and reverse directions through Min Bandwidth Spec and Min Reverse Bandwidth Spec fields.

Optional TLVs MAY be included within the object body to specify more specific bandwidth requirements. No TLVs for the Generalized Load Balancing object type are defined by this document.

The semantic of the LOAD-BALANCING object is not changed. If a PCC requests the computation of a set of TE-LSPs with at most N TE-LSPs so that it can carry Generalized bandwidth X, each TE-LSP must at least transport bandwidth B; it inserts a BANDWIDTH object specifying X as the required bandwidth and a LOAD-BALANCING object with the Max-LSP and Min Bandwidth Spec fields set to N and B, respectively. When the BANDWIDTH and Min Bandwidth Spec can be summarized as scalars, the sum of the bandwidth for all TE-LSPs in the set is greater than X. The mapping of the X over N path with (at least) bandwidth B is technology and possibly node specific. Each standard definition of the transport technology is defining those mappings and are not repeated in this document. A simplified example for SDH is described in Appendix A.

In all other cases, including technologies based on statistical multiplexing (e.g., InterServ and Ethernet), the exact bandwidth management (e.g., the Ethernet's Excessive Rate) is left to the PCE's policies, according to the operator's configuration. If required, further documents may introduce a new mechanism to finely express complex load-balancing policies within PCEP.

The BANDWIDTH and LOAD-BALANCING Bw Spec Type can be different depending on the architecture of the endpoint node. When the PCE is not able to handle those two Bw Spec Types, it MUST return a NO-PATH with the bit "LOAD-BALANCING could not be performed with the bandwidth constraints" set in the NO-PATH-VECTOR TLV.

2.5. END-POINTS Object Extensions

The END-POINTS object is used in a PCEP request message to specify the source and the destination of the path for which a path computation is requested. Per [RFC5440], the source IP address and the destination IP address are used to identify those. A new object type is defined to address the following possibilities:

- * Different source and destination endpoint types.
- * Label restrictions on the endpoint.
- * Specification of unnumbered endpoints type as seen in GMPLS networks.

The object encoding is described in the following sections.

In path computation within a GMPLS context, the endpoints can:

- * Be unnumbered as described in [RFC3477].

Reserved bits SHOULD be set to 0 when a message is sent and ignored when the message is received.

The values for the Endpoint Type field are defined as follows:

Value	Type
0	Point-to-Point
1	Point-to-Multipoint with leaf type 1
2	Point-to-Multipoint with leaf type 2
3	Point-to-Multipoint with leaf type 3
4	Point-to-Multipoint with leaf type 4
5-244	Unassigned
245-255	Experimental Use

Table 4: Generalized Endpoint Types

The Endpoint Type field is used to cover both point-to-point and different point-to-multipoint endpoints. A PCE may only accept endpoint type 0; endpoint types 1-4 apply if the PCE implementation supports P2MP path calculation. The leaf types for P2MP are as per [RFC8306]. A PCE not supporting a given endpoint type SHOULD respond with a PCerr with Error-Type=4 (Not supported object) and Error-value=7 (Unsupported endpoint type in END-POINTS Generalized Endpoint object type). As per [RFC5440], a PCE unable to process Generalized Endpoints may respond with Error-Type=3 (Unknown Object) and Error-value=2 (Unrecognized object type) or with Error-Type=4 (Not supported object) and Error-value=2 (Not supported object Type). The TLVs present in the request object body MUST follow the grammar per [RFC5511]:

```

<generalized-endpoint-tlvs> ::=
  <p2p-endpoints> | <p2mp-endpoints>

<p2p-endpoints> ::=
  <endpoint> [<endpoint-restriction-list>]
  <endpoint> [<endpoint-restriction-list>]

<p2mp-endpoints> ::=
  <endpoint> [<endpoint-restriction-list>]
  <endpoint> [<endpoint-restriction-list>]
  [<endpoint> [<endpoint-restriction-list>]]...

```

For endpoint type Point-to-Point, two endpoint TLVs MUST be present in the message. The first endpoint is the source, and the second is the destination.

For endpoint type Point-to-Multipoint, several END-POINTS objects MAY be present in the message, and the exact meaning depends on the endpoint type defined for the object. The first endpoint TLV is the root, and other endpoint TLVs are the leaves. The root endpoint MUST be the same for all END-POINTS objects for that P2MP tree request. If the root endpoint is not the same for all END-POINTS, a PCerr with Error-Type=17 (P2MP END-POINTS Error) and Error-value=4 (The PCE cannot satisfy the request due to inconsistent END-POINTS) MUST be returned. The procedure defined in [RFC8306], Section 3.10 also applies to the Generalized Endpoint with Point-to-Multipoint endpoint

types.

An endpoint is defined as follows:

```
<endpoint> ::= <IPV4-ADDRESS> | <IPV6-ADDRESS> | <UNNUMBERED-ENDPOINT>
<endpoint-restriction-list> ::= <endpoint-restriction>
                                [<endpoint-restriction-list>]

<endpoint-restriction> ::=
                                [<LABEL-REQUEST>][<label-restriction-list>]

<label-restriction-list> ::= <label-restriction>
                                [<label-restriction-list>]

<label-restriction> ::= <LABEL-SET>
```

The different TLVs are described in the following sections. A PCE MAY support any or all of the IPV4-ADDRESS, IPV6-ADDRESS, and UNNUMBERED-ENDPOINT TLVs. When receiving a PCReq, a PCE unable to resolve the identifier in one of those TLVs MUST respond by using a PCRep with NO-PATH and setting the bit "Unknown destination" or "Unknown source" in the NO-PATH-VECTOR TLV. The response SHOULD include the END-POINTS object with only the unsupported TLV(s).

A PCE MAY support either or both of the LABEL-REQUEST and LABEL-SET TLVs. If a PCE finds a non-supported TLV in the END-POINTS, the PCE MUST respond with a PCErr message with Error-Type=4 (Not supported object) and Error-value=8 (Unsupported TLV present in END-POINTS Generalized Endpoint object type), and the message SHOULD include the END-POINTS object in the response with only the endpoint and endpoint restriction TLV it did not understand. A PCE supporting those TLVs but not being able to fulfill the label restriction MUST send a response with a NO-PATH object that has the bit "No endpoint label resource" or "No endpoint label resource in range" set in the NO-PATH-VECTOR TLV. The response SHOULD include an END-POINTS object containing only the TLV(s) related to the constraints the PCE could not meet.

2.5.2. END-POINTS TLV Extensions

All endpoint TLVs have the standard PCEP TLV header as defined in [RFC5440], Section 7.1. For the Generalized Endpoint object type, the TLVs MUST follow the ordering defined in Section 2.5.1.

2.5.2.1. IPV4-ADDRESS TLV

The IPV4-ADDRESS TLV (Type 39) represents a numbered endpoint using IPv4 numbering. The format of the TLV value is as follows:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

This TLV MAY be ignored, in which case a PCRep with NO-PATH SHOULD be returned, as described in Section 2.5.1.

2.5.2.2. IPV6-ADDRESS TLV

The IPV6-ADDRESS TLV (Type 40) represents a numbered endpoint using IPv6 numbering. The format of the TLV value is as follows:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv6 address (16 bytes)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```



```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Action | Reserved | L | O | U | Label Type |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Subchannel 1 |
| ... |
+-----+-----+-----+-----+-----+-----+-----+-----+
: :
: :
+-----+-----+-----+-----+-----+-----+-----+-----+
| Subchannel N |
| ... |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

A LABEL-SET TLV represents a set of possible labels that can be used on an interface. If the L bit is cleared, the label allocated on the first endpoint MUST be within the label set range. The Action parameter in the LABEL-SET indicates the type of list provided. These parameters are described by [RFC3471], Section 3.5.1.

The U, O, and L bits are defined as follows:

- U: Upstream direction. Set for the upstream (reverse) direction in case of bidirectional LSP.
- O: Old label. Set when the TLV represents the old (previously allocated) label in case of reoptimization. The R bit of the RP object MUST be set to 1. If the L bit is set, this bit SHOULD be set to 0 and ignored on receipt. When this bit is set, the Action field MUST be set to 0 (Inclusive List), and the LABEL-SET MUST contain one subchannel.
- L: Loose label. Set when the TLV indicates to the PCE that a set of preferred (ordered) labels are to be used. The PCE MAY use those labels for label allocation.

Several LABEL_SET TLVs MAY be present with the O bit cleared; LABEL_SET TLVs with the L bit set can be combined with a LABEL_SET TLV with the L bit cleared. There MUST NOT be more than two LABEL_SET TLVs present with the O bit set. If there are two LABEL_SET TLVs present, there MUST NOT be more than one with the U bit set, and there MUST NOT be more than one with the U bit cleared. For a given U bit value, if more than one LABEL_SET TLV with the O bit set is present, the first TLV MUST be processed, and the following TLVs that have the same U and O bits MUST be ignored.

A LABEL-SET TLV with the O and L bits set MUST trigger a PCErr message with Error-Type=10 (Reception of an invalid object) and Error-value=29 (Wrong LABEL-SET TLV present with O and L bits set).

A LABEL-SET TLV that has the O bit set and an Action field not set to 0 (Inclusive List) or that contains more than one subchannel MUST trigger a PCErr message with Error-Type=10 (Reception of an invalid object) and Error-value=30 (Wrong LABEL-SET TLV present with O bit set and wrong format).

If a LABEL-SET TLV is present with the O bit set, the R bit of the RP object MUST be set; otherwise, a PCErr message MUST be sent with Error-Type=10 (Reception of an invalid object) and Error-value=28 (LABEL-SET TLV present with O bit set but without R bit set in RP).

2.6. IRO Extension

The IRO as defined in [RFC5440] is used to include specific objects in the path. RSVP-TE allows the inclusion of a label definition. In order to fulfill requirement 13 in Section 3.1 of [RFC7025], the IRO needs to support the new subobject type as defined in [RFC3473]:

currently an IP address subobject (Type 1 or 2) or an interface ID (Type 4) subobject. If an IP address subobject is used, the given IP address MUST be associated with a link. More than one label subobject MAY follow a subobject identifying a link.

+=====+	
Type	Subobject
+=====+	
10	Label
+-----+	

Table 6

2.8. LSPA Extensions

The LSPA carries the LSP attributes. In the end-to-end recovery context, this also includes the protection state information. A new TLV is defined to fulfill requirement 7 in Section 3.1 of [RFC7025] and requirement 3 in Section 3.2 of [RFC7025]. This TLV contains the information of the PROTECTION object defined by [RFC4872] and can be used as a policy input. The LSPA object MAY carry a PROTECTION-ATTRIBUTE TLV (Type 44), which is defined as follows:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+-----+-----+										+-----+-----+										+-----+-----+										+-----+-----+									
Type										Length																													
+-----+-----+										+-----+-----+										+-----+-----+										+-----+-----+									
S P N O					Reserved					LSP Flags					Reserved					Link Flags																			
+-----+-----+										+-----+-----+										+-----+-----+										+-----+-----+									
I R					Reserved					Seg.Flags					Reserved																								
+-----+-----+										+-----+-----+										+-----+-----+										+-----+-----+									

The content is as defined in [RFC4872], Section 14 and [RFC4873], Section 6.1.

The LSP (protection) Flags field or the Link Flags field can be used by a PCE implementation for routing policy input. The other attributes are only meaningful for a stateful PCE.

This TLV is OPTIONAL and MAY be ignored by the PCE. If ignored by the PCE, it MUST NOT include the TLV in the LSPA of the response. When the TLV is used by the PCE, an LSPA object and the PROTECTION-ATTRIBUTE TLV MUST be included in the response. Fields that were not considered MUST be set to 0.

2.9. NO-PATH Object Extension

The NO-PATH object is used in PCRep messages in response to an unsuccessful Path Computation Request (the PCE could not find a path satisfying the set of constraints). In this scenario, the PCE MUST include a NO-PATH object in the PCRep message. The NO-PATH object MAY carry the NO-PATH-VECTOR TLV that specifies more information on the reasons that led to a negative reply. In case of GMPLS networks, there could be some additional constraints that led to the failure such as protection mismatch, lack of resources, and so on. Several new flags have been defined in the 32-bit Flag field of the NO-PATH-VECTOR TLV, but no modifications have been made in the NO-PATH object.

2.9.1. Extensions to NO-PATH-VECTOR TLV

The modified NO-PATH-VECTOR TLV carrying the additional information is as follows:

Bit number 18: Protection Mismatch (1 bit). Specifies the

mismatch of the protection type in the PROTECTION-ATTRIBUTE TLV in the request.

Bit number 17: No Resource (1 bit). Specifies that the resources are not currently sufficient to provide the path.

Bit number 16: Granularity not supported (1 bit). Specifies that the PCE is not able to provide a path with the requested granularity.

Bit number 15: No endpoint label resource (1 bit). Specifies that the PCE is not able to provide a path because of the endpoint label restriction.

Bit number 14: No endpoint label resource in range (1 bit). Specifies that the PCE is not able to provide a path because of the endpoint label set restriction.

Bit number 13: No label resource in range (1 bit). Specifies that the PCE is not able to provide a path because of the label set restriction.

Bit number 12: LOAD-BALANCING could not be performed with the bandwidth constraints (1 bit). Specifies that the PCE is not able to provide a path because it could not map the BANDWIDTH into the parameters specified by the LOAD-BALANCING.

3. Additional Error-Types and Error-Values Defined

A PCEP-ERROR object is used to report a PCEP error and is characterized by an Error-Type that specifies the type of error and an Error-value that provides additional information about the error. An additional Error-Type and several Error-values are defined to represent some of the errors related to the newly identified objects, which are related to GMPLS networks. For each PCEP error, an Error-Type and an Error-value are defined. Error-Types 1 to 10 are already defined in [RFC5440]. Additional Error-values are defined for Error-Types 4 and 10. A new Error-Type 29 (Path computation failure) is defined in this document.

Error-Type 29 (Path computation failure) is used to reflect constraints not understood by the PCE, for instance, when the PCE is not able to understand the Generalized bandwidth. If the constraints are understood, but the PCE is unable to find those constraints, NO-PATH is to be used.

Error-Type	Meaning	Error-value
4	Not supported object	
		6: BANDWIDTH object type 3 or 4 not supported
		7: Unsupported endpoint type in END-POINTS Generalized Endpoint object type
		8: Unsupported TLV present in END-POINTS Generalized Endpoint object type
		9: Unsupported granularity in the RP object flags

10	Reception of an invalid object	
		24: Bad BANDWIDTH object type 3 or 4
		25: Unsupported LSP Protection Flags in PROTECTION-ATTRIBUTE TLV
		26: Unsupported Secondary LSP Protection Flags in PROTECTION-ATTRIBUTE TLV
		27: Unsupported Link Protection Type in PROTECTION-ATTRIBUTE TLV
		28: LABEL-SET TLV present with O bit set but without R bit set in RP
		29: Wrong LABEL-SET TLV present with O and L bits set
		30: Wrong LABEL-SET TLV present with O bit set and wrong format
		31: Missing GMPLS-CAPABILITY TLV
29	Path computation failure	
		0: Unassigned
		1: Unacceptable request message
		2: Generalized bandwidth value not supported
		3: Label set constraint could not be met
		4: Label constraint could not be met

Table 7

4. Manageability Considerations

This section follows the guidance of [RFC6123].

4.1. Control of Function through Configuration and Policy

This document makes no change to the basic operation of PCEP, so the requirements described in [RFC5440], Section 8.1 also apply to this document. In addition to those requirements, a PCEP implementation may allow the configuration of the following parameters:

- * Accepted RG in the RP object.
- * Default RG to use (overriding the one present in the PCReq).
- * Accepted BANDWIDTH object type 3 and 4 parameters in the request and default mapping to use when not specified in the request.
- * Accepted LOAD-BALANCING object type 2 parameters in request.
- * Accepted endpoint type and allowed TLVs in object END-POINTS with the object type Generalized Endpoint.
- * Accepted range for label restrictions in END-POINTS or IRO/XRO objects.
- * Acceptance and suppression of the PROTECTION-ATTRIBUTE TLV.

The configuration of the above parameters is applicable to the different sessions as described in [RFC5440], Section 8.1 (by default, per PCEP peer, etc.).

4.2. Information and Data Models

This document makes no change to the basic operation of PCEP, so the requirements described in [RFC5440], Section 8.2 also apply to this document. This document does not introduce any new ERO subobjects; the ERO information model is already covered in [RFC4802].

4.3. Liveness Detection and Monitoring

This document makes no change to the basic operation of PCEP, so there are no changes to the requirements for liveness detection and monitoring in [RFC4657] and [RFC5440], Section 8.3.

4.4. Verifying Correct Operation

This document makes no change to the basic operations of PCEP and the considerations described in [RFC5440], Section 8.4. New errors defined by this document should satisfy the requirement to log error events.

4.5. Requirements on Other Protocols and Functional Components

No new requirements on other protocols and functional components are made by this document. This document does not require ERO object extensions. Any new ERO subobject defined in the TEAS or CCAMP Working Groups can be adopted without modifying the operations defined in this document.

4.6. Impact on Network Operation

This document makes no change to the basic operations of PCEP and the considerations described in [RFC5440], Section 8.6. In addition to the limit on the rate of messages sent by a PCEP speaker, a limit MAY be placed on the size of the PCEP messages.

5. IANA Considerations

IANA assigns values to PCEP objects and TLVs. IANA has made allocations for the newly defined objects and TLVs defined in this document. In addition, IANA manages the space of flags that have been newly added in the TLVs.

5.1. PCEP Objects

New object types are defined in Sections 2.3, 2.4, and 2.5.1. IANA

has made the following Object-Type allocations in the "PCEP Objects" subregistry.

Object-Class Value	Name	Object-Type	Reference
5	BANDWIDTH	3: Generalized bandwidth	RFC 8779, Section 2.3
		4: Generalized bandwidth of an existing TE-LSP for which a reoptimization is requested	RFC 8779, Section 2.3
14	LOAD-BALANCING	2: Generalized Load Balancing	RFC 8779, Section 2.4
4	END-POINTS	5: Generalized Endpoint	RFC 8779, Section 2.5

Table 8

5.2. Endpoint Type Field in the Generalized END-POINTS Object

IANA has created a new "Generalized Endpoint Types" registry to manage the Endpoint Type field of the END-POINTS object, the object type Generalized Endpoint, and the code space.

New endpoint types in the Unassigned range are assigned by Standards Action [RFC8126]. Each endpoint type should be tracked with the following attributes:

- * Value
- * Type
- * Defining RFC

New endpoint types in the Experimental Use range will not be registered with IANA and MUST NOT be mentioned by any RFCs.

The following values are defined by this document (see Table 4 in Section 2.5.1):

Value	Type
0	Point-to-Point
1	Point-to-Multipoint with leaf type 1
2	Point-to-Multipoint with leaf type 2
3	Point-to-Multipoint with leaf type 3
4	Point-to-Multipoint with leaf type 4
5-244	Unassigned

245-255 Experimental Use
+-----+

Table 9

5.3. New PCEP TLVs

IANA manages a registry for PCEP TLV code points (see [RFC5440]), which is maintained as the "PCEP TLV Type Indicators" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry. IANA has allocated the following per this document:

Value	Meaning	Reference
39	IPV4-ADDRESS	RFC 8779, Section 2.5.2.1
40	IPV6-ADDRESS	RFC 8779, Section 2.5.2.2
41	UNNUMBERED-ENDPOINT	RFC 8779, Section 2.5.2.3
42	LABEL-REQUEST	RFC 8779, Section 2.5.2.4
43	LABEL-SET	RFC 8779, Section 2.5.2.5
44	PROTECTION-ATTRIBUTE	RFC 8779, Section 2.8
45	GMPLS-CAPABILITY	RFC 8779, Section 2.1.2

Table 10

5.4. RP Object Flag Field

A new flag is defined in Section 2.2 for the Flags field of the RP object. IANA has made the following allocation in the "RP Object Flag Field" subregistry:

Bit	Description	Reference
15-16	Routing Granularity (RG)	RFC 8779, Section 2.2

Table 11

5.5. New PCEP Error Codes

New PCEP Error-Types and Error-values are defined in Section 3. IANA has made the following allocations in the "PCEP-ERROR Object Error Types and Values" registry:

Error-Type	Meaning	Error-value	Reference
4	Not supported object		[RFC5440]
		6: BANDWIDTH object type 3 or 4 not supported	RFC 8779
		7: Unsupported endpoint type in END-POINTS Generalized Endpoint object type	RFC 8779

		8: Unsupported TLV present in END-POINTS Generalized Endpoint object type	RFC 8779
		9: Unsupported granularity in the RP object flags	RFC 8779
10	Reception of an invalid object		[RFC5440]
		24: Bad BANDWIDTH object type 3 or 4	RFC 8779
		25: Unsupported LSP Protection Flags in PROTECTION-ATTRIBUTE TLV	RFC 8779
		26: Unsupported Secondary LSP Protection Flags in PROTECTION-ATTRIBUTE TLV	RFC 8779
		27: Unsupported Link Protection Type in PROTECTION-ATTRIBUTE TLV	RFC 8779
		28: LABEL-SET TLV present with O bit set but without R bit set in RP	RFC 8779
		29: Wrong LABEL-SET TLV present with O and L bits set	RFC 8779
		30: Wrong LABEL-SET TLV present with O bit set and wrong format	RFC 8779
		31: Missing GMPLS-CAPABILITY TLV	RFC 8779
29	Path computation failure		RFC 8779
		0: Unassigned	RFC 8779
		1: Unacceptable request message	RFC 8779
		2: Generalized bandwidth value not supported	RFC 8779
		3: Label set constraint could not be met	RFC 8779
		4: Label constraint could not be met	RFC 8779

Table 12

5.6. New Bits in NO-PATH-VECTOR TLV

New NO-PATH-VECTOR TLV bits are defined in Section 2.9.1. IANA has made the following allocations in the "NO-PATH-VECTOR TLV Flag Field" subregistry:

Bit	Description	Reference
18	Protection Mismatch	RFC 8779
17	No Resource	RFC 8779
16	Granularity not supported	RFC 8779
15	No endpoint label resource	RFC 8779
14	No endpoint label resource in range	RFC 8779
13	No label resource in range	RFC 8779
12	LOAD-BALANCING could not be performed with the bandwidth constraints	RFC 8779

Table 13

5.7. New Subobject for the Include Route Object

IANA has added a new subobject in the "IRO Subobjects" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry.

IANA has added a new subobject that can be carried in the IRO as follows:

Value	Description	Reference
10	Label	RFC 8779

Table 14

5.8. New Subobject for the Exclude Route Object

IANA has added a new subobject in the "XRO Subobjects" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry.

IANA has added a new subobject that can be carried in the XRO as follows:

Value	Description	Reference
10	Label	RFC 8779

Table 15

5.9. New GMPLS-CAPABILITY TLV Flag Field

IANA has created a new "GMPLS-CAPABILITY TLV Flag Field" subregistry within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flag field of the GMPLS-CAPABILITY TLV.

New bit numbers are to be assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit number (counting from bit 0 as the most significant bit)
- * Capability description
- * Defining RFC

The initial contents of the subregistry are empty, with bits 0-31 marked as Unassigned.

6. Security Considerations

GMPLS controls multiple technologies and types of network elements. The LSPs that are established using GMPLS, whose paths can be computed using the PCEP extensions to support GMPLS described in this document, can carry a high volume of traffic and can be a critical part of a network infrastructure. The PCE can then play a key role in the use of the resources and in determining the physical paths of the LSPs; thus, it is important to ensure the identity of the PCE and PCC, as well as the communication channel. In many deployments, there will be a completely isolated network where an external attack is of very low probability. However, there are other deployment cases in which the PCC-PCE communication can be more exposed, and there could be more security considerations. There are three main situations in case an attack in the GMPLS PCE context happens:

PCE Identity theft: A legitimate PCC could request a path for a GMPLS LSP to a malicious PCE, which poses as a legitimate PCE. The response may be that the LSP traverses some geographical place known to the attacker where confidentiality (sniffing), integrity (traffic modification), or availability (traffic drop) attacks could be performed by use of an attacker-controlled middlebox device. Also, the resulting LSP can omit constraints given in the requests (e.g., excluding certain fibers and avoiding some SRLGs), which could make the LSP that will be set up later look perfectly fine, but it will be in a risky situation. Also, the result can lead to the creation of an LSP that does not provide the desired quality and gives less resources than necessary.

PCC Identity theft: A malicious PCC, acting as a legitimate PCC, requesting LSP paths to a legitimate PCE can obtain a good knowledge of the physical topology of a critical infrastructure. It could learn enough details to plan a later physical attack.

Message inspection: As in the previous case, knowledge of an infrastructure can be obtained by sniffing PCEP messages.

The security mechanisms can provide authentication and confidentiality for those scenarios where PCC-PCE communication cannot be completely trusted. [RFC8253] provides origin verification, message integrity, and replay protection, and it ensures that a third party cannot decipher the contents of a message.

In order to protect against the malicious PCE case, the PCC SHOULD have policies in place to accept or not accept the path provided by the PCE. Those policies can verify if the path follows the provided constraints. In addition, a technology-specific data-plane mechanism can be used (following [RFC5920], Section 5.8) to verify the data-plane connectivity and deviation from constraints.

The usage of Transport Layer Security (TLS) to enhance PCEP security

is described in [RFC8253]. The document describes the initiation of TLS procedures, the TLS handshake mechanisms, the TLS methods for peer authentication, the applicable TLS ciphersuites for data exchange, and the handling of errors in the security checks. PCE and PCC SHOULD use the mechanism in [RFC8253] to protect against malicious PCC and PCE.

Finally, as mentioned by [RFC7025], the PCEP extensions that support GMPLS should be considered under the same security as current PCE work, and this extension will not change the underlying security issues. However, given the critical nature of the network infrastructures under control by GMPLS, the security issues described above should be seriously considered when deploying a GMPLS-PCE-based control plane for such networks. For an overview of the security considerations, not only related to PCE/PCEP, and vulnerabilities of a GMPLS control plane, see [RFC5920].

7. References

7.1. Normative References

- [G.709-v3] ITU-T, "Interfaces for the optical transport network", Recommendation G.709/Y.1331, June 2016, <<https://www.itu.int/rec/T-REC-G.709-201606-I/en>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, DOI 10.17487/RFC2210, September 1997, <<https://www.rfc-editor.org/info/rfc2210>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, DOI 10.17487/RFC3471, January 2003, <<https://www.rfc-editor.org/info/rfc3471>>.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, DOI 10.17487/RFC3473, January 2003, <<https://www.rfc-editor.org/info/rfc3473>>.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, DOI 10.17487/RFC3477, January 2003, <<https://www.rfc-editor.org/info/rfc3477>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC4003] Berger, L., "GMPLS Signaling Procedure for Egress Control", RFC 4003, DOI 10.17487/RFC4003, February 2005, <<https://www.rfc-editor.org/info/rfc4003>>.
- [RFC4328] Papadimitriou, D., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for G.709 Optical

- Transport Networks Control", RFC 4328,
DOI 10.17487/RFC4328, January 2006,
<<https://www.rfc-editor.org/info/rfc4328>>.
- [RFC4606] Mannie, E. and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 4606,
DOI 10.17487/RFC4606, August 2006,
<<https://www.rfc-editor.org/info/rfc4606>>.
- [RFC4802] Nadeau, T., Ed. and A. Farrel, Ed., "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", RFC 4802,
DOI 10.17487/RFC4802, February 2007,
<<https://www.rfc-editor.org/info/rfc4802>>.
- [RFC4872] Lang, J.P., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, DOI 10.17487/RFC4872, May 2007,
<<https://www.rfc-editor.org/info/rfc4872>>.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, DOI 10.17487/RFC4873, May 2007, <<https://www.rfc-editor.org/info/rfc4873>>.
- [RFC5088] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, DOI 10.17487/RFC5088, January 2008, <<https://www.rfc-editor.org/info/rfc5088>>.
- [RFC5089] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, DOI 10.17487/RFC5089, January 2008, <<https://www.rfc-editor.org/info/rfc5089>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009,
<<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, DOI 10.17487/RFC5511, April 2009, <<https://www.rfc-editor.org/info/rfc5511>>.
- [RFC5520] Bradford, R., Ed., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, DOI 10.17487/RFC5520, April 2009,
<<https://www.rfc-editor.org/info/rfc5520>>.
- [RFC5521] Oki, E., Takeda, T., and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", RFC 5521, DOI 10.17487/RFC5521, April 2009, <<https://www.rfc-editor.org/info/rfc5521>>.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", RFC 5541, DOI 10.17487/RFC5541, June 2009,
<<https://www.rfc-editor.org/info/rfc5541>>.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiimoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol

Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, DOI 10.17487/RFC6001, October 2010, <<https://www.rfc-editor.org/info/rfc6001>>.

- [RFC6003] Papadimitriou, D., "Ethernet Traffic Parameters", RFC 6003, DOI 10.17487/RFC6003, October 2010, <<https://www.rfc-editor.org/info/rfc6003>>.
- [RFC6205] Otani, T., Ed. and D. Li, Ed., "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, DOI 10.17487/RFC6205, March 2011, <<https://www.rfc-editor.org/info/rfc6205>>.
- [RFC6387] Takacs, A., Berger, L., Caviglia, D., Fedyk, D., and J. Meuric, "GMPLS Asymmetric Bandwidth Bidirectional Label Switched Paths (LSPs)", RFC 6387, DOI 10.17487/RFC6387, September 2011, <<https://www.rfc-editor.org/info/rfc6387>>.
- [RFC7139] Zhang, F., Ed., Zhang, G., Belotti, S., Ceccarelli, D., and K. Pithewan, "GMPLS Signaling Extensions for Control of Evolving G.709 Optical Transport Networks", RFC 7139, DOI 10.17487/RFC7139, March 2014, <<https://www.rfc-editor.org/info/rfc7139>>.
- [RFC7570] Margaria, C., Ed., Martinelli, G., Balls, S., and B. Wright, "Label Switched Path (LSP) Attribute in the Explicit Route Object (ERO)", RFC 7570, DOI 10.17487/RFC7570, July 2015, <<https://www.rfc-editor.org/info/rfc7570>>.
- [RFC7792] Zhang, F., Zhang, X., Farrel, A., Gonzalez de Dios, O., and D. Ceccarelli, "RSVP-TE Signaling Extensions in Support of Flexi-Grid Dense Wavelength Division Multiplexing (DWDM) Networks", RFC 7792, DOI 10.17487/RFC7792, March 2016, <<https://www.rfc-editor.org/info/rfc7792>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <<https://www.rfc-editor.org/info/rfc8253>>.
- [RFC8282] Oki, E., Takeda, T., Farrel, A., and F. Zhang, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 8282, DOI 10.17487/RFC8282, December 2017, <<https://www.rfc-editor.org/info/rfc8282>>.
- [RFC8306] Zhao, Q., Dhody, D., Ed., Palleti, R., and D. King, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", RFC 8306, DOI 10.17487/RFC8306, November 2017, <<https://www.rfc-editor.org/info/rfc8306>>.

7.2. Informative References

- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.
- [RFC4657] Ash, J., Ed. and J.L. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, DOI 10.17487/RFC4657, September 2006, <<https://www.rfc-editor.org/info/rfc4657>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.
- [RFC6123] Farrel, A., "Inclusion of Manageability Sections in Path Computation Element (PCE) Working Group Drafts", RFC 6123, DOI 10.17487/RFC6123, February 2011, <<https://www.rfc-editor.org/info/rfc6123>>.
- [RFC6163] Lee, Y., Ed., Bernstein, G., Ed., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, DOI 10.17487/RFC6163, April 2011, <<https://www.rfc-editor.org/info/rfc6163>>.
- [RFC7025] Otani, T., Ogaki, K., Caviglia, D., Zhang, F., and C. Margaria, "Requirements for GMPLS Applications of PCE", RFC 7025, DOI 10.17487/RFC7025, September 2013, <<https://www.rfc-editor.org/info/rfc7025>>.
- [RFC7449] Lee, Y., Ed., Bernstein, G., Ed., Martensson, J., Takeda, T., Tsuritani, T., and O. Gonzalez de Dios, "Path Computation Element Communication Protocol (PCEP) Requirements for Wavelength Switched Optical Network (WSON) Routing and Wavelength Assignment", RFC 7449, DOI 10.17487/RFC7449, February 2015, <<https://www.rfc-editor.org/info/rfc7449>>.

Appendix A. LOAD-BALANCING Usage for SDH Virtual Concatenation

As an example, a request for one co-signaled $n \times$ VC-4 TE-LSP will not use LOAD-BALANCING. In case the VC-4 components can use different paths, the BANDWIDTH with object type 3 will contain the complete $n \times$ VC-4 traffic specification, and the LOAD-BALANCING object will contain the minimum co-signaled VC-4. For an SDH network, a request for a TE-LSP group with 10 VC-4 containers, with each path using at minimum 2 \times VC-4 containers, can be represented with a BANDWIDTH object with object type 3, the Bw Spec Type set to 4, and the content of the Generalized Bandwidth field with ST=6, RCC=0, NCC=0, NVC=10, and MT=1. The LOAD-BALANCING with object type 2 with the Bw Spec Type set to 4 and Max-LSP=5, Min Bandwidth Spec is ST=6, RCC=0, NCC=0, NVC=2, MT=1. The PCE can respond with a maximum of 5 paths, with each path having a BANDWIDTH object type 3 and a Generalized Bandwidth field matching the Min Bandwidth Spec from the LOAD-BALANCING object of the corresponding request.

Acknowledgments

The research of Ramon Casellas, Francisco Javier Jimenez Chico, Oscar Gonzalez de Dios, Cyril Margaria, and Franz Rambach that led to the results in this document received funding from the European Community's Seventh Framework Program FP7/2007-2013 under grant agreement no. 247674 and no. 317999.

The authors would like to thank Julien Meuric, Lyndon Ong, Giada

Lander, Jonathan Hardwick, Diego Lopez, David Sinicrope, Vincent Roca, Dhruv Dhody, Adrian Farrel, and Tianran Zhou for their review and useful comments.

Thanks to Alisa Cooper, Benjamin Kaduk, Elwyn Davies, Martin Vigoureux, Roman Danyliw, and Suresh Krishnan for the IESG-related comments.

Contributors

Elie Sfeir
Coriant
St. Martin Strasse 76
81541 Munich
Germany

Email: elie.sfeir@coriant.com

Franz Rambach
Nockherstrasse 2-4
81541 Munich
Germany

Phone: +49 178 8855738
Email: franz.rambach@cgi.com

Francisco Javier Jimenez Chico
Telefonica Investigacion y Desarrollo
C/ Emilio Vargas 6
28043 Madrid
Spain

Phone: +34 91 3379037
Email: fjjc@tid.es

Suresh Babu

Email: sureshhimnish@gmail.com

Young Lee
Samsung Electronics

Email: younglee.tx@gmail.com

Senthil Kumar S

Email: ssenthilkumar@gmail.com

Jun Sun
Huawei Technologies
Shenzhen
China

Email: johnsun@huawei.com

Ramon Casellas
CTTC - Centre Tecnologic de Telecomunicacions de Catalunya
PMT Ed B4 Av. Carl Friedrich Gauss 7
08660 Castelldefels, Barcelona

Spain

Phone: +34 93 6452916

Email: ramon.casellas@cttc.e

Authors' Addresses

Cyril Margaria (editor)

Juniper

Email: cmargaria@juniper.net

Oscar Gonzalez de Dios (editor)

Telefonica Investigacion y Desarrollo

C/ Ronda de la Comunicacion

28050 Madrid

Spain

Phone: +34 91 4833441

Email: oscar.gonzalezdedios@telefonica.com

Fatai Zhang (editor)

Huawei Technologies

F3-5-B R&D Center, Huawei Base

Bantian, Longgang District

Shenzhen

518129

China

Email: zhangfatai@huawei.com